REMARKS

Claims 1-27 are pending in the present application. Claims 1, 2, 5-8, 11, 12, 14-19, 21-25, and 27 are amended. Reconsideration of the claims is respectfully requested.

Amendments were made to the specification to correct errors and to clarify the specification. No new matter has been added by any of the amendments to the specification.

I. 35 U.S.C. § 102, Anticipation

The Office Action rejects claims 1-27 under 35 U.S.C. § 102 as being anticipated by *Chong, Jr.* (US Patent No. 6,349,357). This rejection is respectfully traversed.

As to claims 1 and 7, the Office Action states:

> With respect to claims 1 and 7, Chong, Jr. discloses a method for managing a read request including receiving, at a first controller, a read request for a data block [see controller 26 in Figure 4A, as well as corresponding parts of Fig. 5, and column 4, lines 68-71], and allocating a memory buffer for the data block from a memory pool that includes a first memory on the first controller and a second memory on a second controller, wherein the memory buffer resides in the second memory [Chong, Jr. teaches that a cache memory may be allocated to buffer data, where the buffer may be allocated from a memory pool including a first cache memory 341 and a second cache memory 342. Chong, Jr. also teaches that space may be allocated in the first cache memory and in the second cache memory to mirror the data and provide improved reliability (see column 7, lines 11-13; column 8, line 63 to column 9, line 1; column 11, lines 64-66; and column 12, lines 6-9 and 17-21, e.g.).].
>
> Also with respect to claim 1, Chong, Jr. further teaches retrieving the data block from a storage device [16 in Fig. 4A and 161, 162 in Figure 5] and caching the data block in the memory buffer [see column 7, lines 11-13, e.g.].

Office Action, dated December 1, 2003. Applicant respectfully disagrees. *Chong, Jr.* teaches a storage architecture that provides scalable performance through independent control and data transfer paths. Commands and status information are passed to a control module in a storage controller through a control transfer path. Data may be passed from a host computer to a storage device and from the storage device to the host computer directly through a data path. The storage controller includes a switch for routing control

and data based on the messaging scheme. See *Chong, Jr.*, Abstract; col. 2, line 56, to col. 4, line 5.

Chong, Jr. also teaches a fault-tolerant configuration with scalable performance storage architecture. This architecture includes switches with additional ports to allow two hosts to simultaneously access each switch for pertinent data transfer operations involving a storage device. Each switch may send data to multiple places at the same time, such as mirrored cache memories. The described scalable performance storage architecture is shown in **Figure 5**, which is reproduced below:

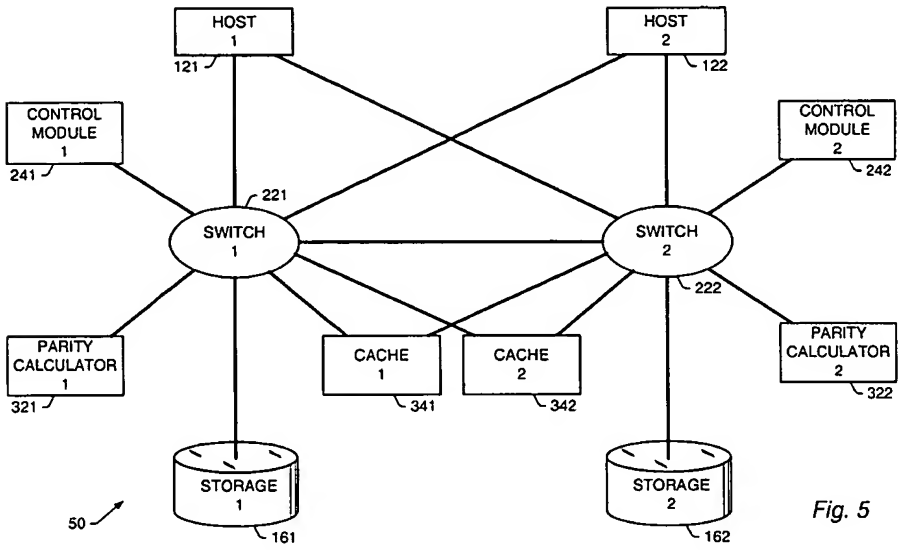

Fig. 5

This architecture appears to be a configuration of separate parts, rather than discrete storage controllers. In other words, switch 1 **221** is not in or a part of a first storage controller and switch 2 **222** is not part of a second storage controller. Rather, the configuration of elements, when taken together, appears to form the architecture.

In contradistinction, the present invention provides a storage controller that allocates a memory buffer for a data block from a memory pool that includes a first memory on the first storage controller and a second memory on a second storage controller and allows the first storage controller to access cached data blocks stored on the second storage controller. Independent claim 1 recites:

> 1. A method for managing a read request, comprising:
>
> receiving, at a first storage controller, a read request for a data block;
>
> allocating a memory buffer for the data block from a memory pool that includes a first memory **on the first storage controller** and a second

memory **on a second storage controller**, wherein the memory buffer resides in the second memory;

retrieving the data block from a storage device; and

caching the data block in the memory buffer.

Independent claim 7 recites:

7. A method for managing a read request, comprising:

receiving, at a first storage controller, a read request for a data block; and

retrieving the data block from a memory pool that includes a first memory on the first storage controller and **a second memory on a second storage controller, wherein the data block resides in the second memory**.

While *Chong, Jr.* does teach an architecture in which a host may access one of two of cache memories through one or two switches, *Chong, Jr.* does not teach or suggest a storage controller that is capable of allocating memory on a separate storage controller, caching a data block on a separate storage controller, and/or retrieving a cached data block stored on a separate storage controller, as recited in claims 1 and 7. That is, in the embodiments in *Chong, Jr.* where separate storage controllers are shown, there is no switch-to-switch path. See *Chong, Jr.*, **Figs. 3A-3E** and **4A-4B**. In the embodiment of *Chong, Jr.* where multiple switches are shown with a switch-to-switch path, there are no separate and distinct controllers.

The applied reference does not teach or suggest each and every claim limitation; therefore, claims 1 and 7 are not anticipated by *Chong, Jr.* Since claims 2-5 and 8-10 depend from claims 1 and 7, the same distinctions between *Chong, Jr.* and the invention recited in claims 1 and 7 apply for these claims. Additionally, claims 2-5 and 8-10 recite other additional combinations of features not suggested by the reference. Claims 18-22 recite subject matter addressed above with respect to claims 1-5 and 7-10 and are allowable for the same reasons.

More particularly, as to claims 2, 3, 8, 9, 12, and 13, the Office Action states:

With respect to claims 2, 8 and 12, Chong, Jr. discloses that the first controller includes a first switch (221 in Fig. 5) and the second controller includes a second switch (222 in Fig. 5).

With respect to claims 3, 9 and 13, Chong, Jr. discloses that the first switch and the second switch are coupled using a switch-to-switch path [note the bus or path between the switches in Fig. 5, for example].

Office Action, dated December 1, 2003. Applicant respectfully disagrees. While *Chong, Jr.* does teach a first and second switch and a switch-to-switch path, *Chong, Jr.* does not teach or fairly suggest that a first switch in a first storage controller and a second switch in a second storage controller wherein the switches of the first storage controller and the second storage controller are coupled using a switch-to-switch path. The applied reference does not teach or suggest each and every claim limitation; therefore, claims 2, 3, 8, 9, 12, and 13 are not anticipated by *Chong, Jr.*

As to claim 11, the Office Action states:

> With respect to claim 11, Chong, Jr. also discloses that the system may be used to manage write requests in addition to read requests wherein a write request for a data block is received at the first controller (see column 4, lines 58-59, e.g.). Chong, Jr. discloses that a "primary" data buffer for the data block may be allocated in a first cache memory and a mirror data buffer allocated for mirroring the data block in a second memory, wherein the first memory resides on one of the first controller and a second controller and the second memory resides on the other of the first controller [see column 7, lines 6-11 and column 11, lines 64-6]. Write data for a data block may be stored in the "primary" data buffer and mirrored in the secondary or mirror data buffer as discussed above.

Office Action, dated December 1, 2003. Applicant respectfully disagrees. The present invention provides a storage controller that allocates a memory buffer for a data block from a memory pool that includes a first memory on the first storage controller and a second memory on a second storage controller and allows the first storage controller to access cached data blocks stored on the second storage controller. Independent claim 11 recites:

> 11. A method for managing a write request, comprising:
> receiving, at a first storage controller, a write request for a data block;
> allocating a primary data buffer for the data block in a first memory and a mirror data buffer for the data block in a second memory, wherein the first memory resides on one of the first storage controller and a second storage controller and the second memory resides on the other of the first storage controller and the second storage controller;
> storing write data for the data block in the primary data buffer; and
> mirroring the write data in the mirror data buffer.

While *Chong, Jr.* does teach an architecture in which a host may mirror data blocks in two cache memories through one or two switches, *Chong, Jr.* does not teach or suggest a storage controller that is capable of allocating memory on a separate storage controller, and mirroring a cached data block between the storage controller and a separate storage controller, as recited in claim 11. That is, in the embodiments in *Chong, Jr.* where separate storage controllers are shown, there is no switch-to-switch path. See *Chong, Jr.*, **Figs. 3A-3E** and **4A-4B**. In the embodiment of *Chong, Jr.* where multiple switches are shown with a switch-to-switch path, there are no separate and distinct controllers. See *Chong, Jr.*, **Figure 5**.

The applied reference does not teach or suggest each and every claim limitation; therefore, claim 11 is not anticipated by *Chong, Jr.* Since claims 1217 depend from claim 11, the same distinctions between *Chong, Jr.* and the invention recited in claim 11 apply for these claims. Additionally, claims 1217 recite other additional combinations of features not suggested by the reference. Claims 18 and 23-27 recite subject matter addressed above with respect to claims 11-17 and are allowable for the same reasons.

More particularly, with respect to claim 18, the Office Action states:

> With respect to claim 18, Chong, Jr. discloses an "apparatus" in a first controller [26 in Fig. 4A, e.g., as well as corresponding parts of Fig. 5] including a host adapter that provides a connection to a host [see hosts 121 and 122 in Fig. 5 and also see column 9, lines 8-10], a processor [see CPUs 241, 242 in Fig. 3A and column 5, lines 5-7, e.g.], a memory controller that manages a connection to a memory [see column 7 lines 37-39], a drive adapter that provides a connection to a storage device [the storage devices 16 and 161, 162 include some kind of "adapter" or controller (not shown) for connection to the bus or switch (see column 5, lines 46-49, as well as column 9, lines 8-10, e.g.)], a first switch that connects the host adapter, the processor, the memory controller, and the drive adapter [see switch 221 in Figure 5, e.g.] and a switch-to-switch path that connects the first switch to a second switch on a second controller [see second switch (222 in Fig. 5) coupled to the first switch by a "switch-to-switch path" (note the bus or path between the switches in Fig. 5, for example)].

Office Action, dated December 1, 2003. Applicant respectfully disagrees. While *Chong, Jr.* does teach a first and second switch and a switch-to-switch path, *Chong, Jr.* does not teach or fairly suggest that a first switch in a first storage controller and a second switch in a second storage controller wherein the switches of the first storage controller and the

second storage controller are coupled using a switch-to-switch path. The applied reference does not teach or suggest each and every claim limitation; therefore, claim 18 is not anticipated by *Chong, Jr.* Since claims 19-27 depend from claim 18, the same distinctions between *Chong, Jr.* and the invention recited in claim 18 apply for these claims. Additionally, claims 19-27 recite other additional combinations of features not suggested by the reference, as discussed above.

Therefore, Applicant respectfully requests withdrawal of the rejection of claims 1-10 under 35 U.S.C. § 102.

Furthermore, *Chong, Jr.* does not teach, suggest, or give any incentive to make the needed changes to reach the presently claimed invention. *Chong, Jr.* actually teaches away from the presently claimed invention because it teaches an architecture without separate storage controllers, as opposed to separate storage controllers where the switches are coupled by a switch-to-switch path, as in the presently claimed invention. Absent the Office Action pointing out some teaching or incentive to implement the architecture of *Chong, Jr.* with separate and distinct storage controllers, one of ordinary skill in the art would not be led to modify *Chong, Jr.* to reach the present invention when the reference is examined as a whole. Absent some teaching, suggestion, or incentive to modify *Chong, Jr.* in this manner, the presently claimed invention can be reached only through an improper use of hindsight using Applicant's disclosure as a template to make the necessary changes to reach the claimed invention.

II. Conclusion

It is respectfully urged that the subject application is patentable over the prior art of record and is now in condition for allowance.

The Examiner is invited to call the undersigned at the below-listed telephone number if in the opinion of the Examiner such a telephone conference would expedite or aid the prosecution and examination of this application.

DATE: March 1, 2004

Respectfully submitted,

Stephen R. Tkacs
Reg. No. 46,430
Carstens, Yee & Cahoon, LLP
P.O. Box 802334
Dallas, TX 75380
(972) 367-2001
Agent for Applicants